

WHAT IS CLAIMED IS:

1. A method for determining at least one best-performing content server in response to a request in a network including a plurality of content servers, at least one redirection server, and a plurality of clients, the method comprising the steps of:

creating a plurality of client clusters, wherein each of said plurality of client clusters includes one or more clients having similar network distance properties;

identifying said at least one best-performing content server for each of said plurality of client clusters by determining network distances between each of said plurality of client clusters and each of said plurality of content servers and selecting at least one content server for each of said plurality of client clusters having a minimum network distance there-between; and

mapping each of said plurality of client clusters to a corresponding said at least one identified best-performing content server.

2. The method of claim 1, wherein the redirection server is an authoritative domain name (DN) server which receives said requests from a plurality of local DNS servers.

3. The method of claim 2, wherein the request is a domain name (DN) request.

4. The method of claim 1, further including the step of assigning a selection probability to each of said at least one identified best-performing content server, wherein said selection probability ensures that a maximum service capacity of each of said at least one identified best-

performing content server is never exceeded.

5. The method of claim 1, wherein the step of creating said plurality of client clusters includes the step of obtaining distance and load information from each of said plurality of content servers.

6. The method of claim 5, wherein the step of obtaining distance information includes the steps of:

collecting at said plurality of content servers a plurality of distance tuples wherein each distance tuple comprises one or more of the following: a network distance, a content server identifier, a time-stamp, and a client internet protocol (IP) address; and

collecting at said plurality of content servers a plurality of load tuples wherein each load tuple comprises one or more of the following: a time-stamp, a content-server ID, a client IP address, a number of hits, and a domain index.

7. The method of claim 5, further comprising the step of pulling said plurality of distance and load tuples from each of said plurality of content servers at successive data acquisition intervals and storing the plurality of distance and load tuples at said at least one redirection server.

8. The method of claim 6, further comprising the step of multiplying certain data values of said plurality of distance and load tuples by a weighting factor in each of said successive data acquisition intervals.

9. The method of claim 6, wherein said network distance is computed as one of a round trip time, jitter, bandwidth and packet loss.

10. The method of claim 9, wherein said round trip time is computed by monitoring all data packets transmitted and received by one of said plurality of content servers.

11. The method of claim 1, wherein the step of creating said plurality of client clusters further comprises the steps of:

grouping distance tuples according to content server ID and CIDR prefix similarity to obtain data values for each grouping;

storing said data values at respective leaf nodes of a hierarchical tree structure wherein said hierarchical tree structure includes connected nodes including a root node and a plurality of interior and leaf nodes such that an entire CIDR address space is represented by said root node; and

pruning said hierarchical tree structure to determine said plurality of client clusters.

12. The method of claim 11, wherein said data values stored at each of said respective leaf nodes include for at least one of said plurality of content servers; a sum of network distances, a sum of the squares of network distances, and a total number of received tuples.

13. The method of claim 11, wherein the pruning step further comprises the steps of:
comparing sibling leaf nodes in said tree structure to determine whether there is sufficient similarity;

combining said sibling leaf nodes into a parent node if sufficient similarity exists;
repeating the comparing and combining steps if sufficient similarity exists; and
identifying remaining leaf nodes as said client clusters if sufficient similarity does not exist.

14. The method of claim 1, wherein the step of mapping each of said plurality of client clusters with the corresponding said at least one identified best-performing content server comprises the step of assigning an assignment probability to each of said at least one identified best-performing content server.

15. The method of claim 13, wherein said assignment probabilities for each of said at least one identified best-performing content server is obtained from a flow map characterizing data flow in the network.

16. A system for determining at least one best-performing content server in response to a client request in a network including a plurality of content servers, at least one redirection server, and a plurality of clients, the system comprising:

means for creating a plurality of client clusters, wherein each of said plurality of client clusters includes one or more clients having IP addresses having similar network distance properties;

means for identifying said at least one best-performing content server for each of said plurality of client clusters by determining network distances between each of said plurality of client clusters and each of said plurality of content servers and selecting at least one content server for each of said plurality of client clusters having a minimum network distance therebetween; and

means for mapping each of said plurality of client clusters with a corresponding said at least one identified best-performing content server.

17. The system of claim 16, further including means for assigning a selection probability to each of said at least one identified best-performing content server, wherein said selection probability ensures that a maximum service capacity of each of said at least one identified best-performing content server is never exceeded.

18. The system of claim 16, wherein said means for creating further comprises means for creating said plurality of client clusters further includes means for obtaining distance and load information from each of said plurality of content servers.

19. The system of claim 18, wherein said means for obtaining distance and load information further includes:

means for collecting at said plurality of content servers a plurality of distance tuples wherein each distance tuple comprises one or more of the following: a network distance, a content server identifier, a time-stamp, and a client IP address; and

means for collecting at said plurality of content servers a plurality of load tuples wherein each load tuple comprises one or more of the following: a time-stamp, content-server ID, a client IP address, a number of hits, and a domain index.

20. The system of claim 16, wherein said identifying means further includes:

means for modifying said determined network distance according to the number of distance tuples received.

21. The system of claim 16, wherein said identifying means further includes:

means for maintaining said identification of said at least one identified best-performing content server for an amount of time determined by a confidence level calculation.